Validation of Matching

Ya Le, Eric Bax, Nicola Barbieri, David Garcia Soriano, Jitesh Mehta, and James Li

Abstract—We introduce a technique to compute probably approximately correct (PAC) bounds on precision and recall for matching algorithms. The bounds require some verified matches, but those matches may be used to develop the algorithms. The bounds can be applied to network reconciliation or entity resolution algorithms, which identify nodes in different networks or values in a data set that correspond to the same entity. For network reconciliation, the bounds do not require knowledge of the network generation process.

arXiv:1411.0023v2 [cs.LG] 11 Apr 2016

1 Introduction

Our interest in validation of matching stems from work on large-scale network reconciliation. Network reconciliation is the process of matching nodes in different networks that correspond to the same entity [1], [2], [3]. Network reconciliation can combine information from multiple networks to form a combined network with richer information about the entities, represented by nodes, and their interactions with each other, represented by edges. For example, in social networks, reconciling profiles across networks can produce richer information about people and their relationships. This can be a nontrivial challenge for internet-scale data, because, for example, there are over 35,000 people named John Smith on LinkedIn. In biology, matching proteins across species based on their webs of interactions with other proteins can help us induce which proteins play similar roles in different organisms. By examining which structures are preserved across different proteins that have similar roles, we can learn which structures are fundamental to the function of those proteins.

Some network reconciliation algorithms are iterative. They establish some initial "seed matches" based on node data or by matching the highest-degree nodes in the different networks to each other [2], [4]. Then they use shared connections to establish more matches, based on the idea that if more of the connections to a node in one network match connections to a node in another network, then the nodes are more likely to match. With each iteration, more matches may be established, offering opportunities for further matches. This iterative matching process is called percolation. Refer to [5] for some theory on how many seeds are needed for percolation to succeed.

Some network reconciliation algorithms do not require a seed set. Nunes et. al. [6] reduce the problem to a binary classification task, where the goal is to classify pairs of nodes as either matched or not matched. Zaslavskiy et. al. [3] propose a convex-concave programming approach by rewriting the problem as a least squares problem on the set of permutation matrices. Pedarsani et. al. [7] cast the problem in a Bayesian framework based on a definition of the probability of correctly mapping two nodes. Refer to [8] for a review of network reconciliation algorithms.

Our matching problem setting is similar to the transductive setting for classification, from Vapnik [9], where there is a set of training examples with known inputs and class labels and a set of working examples with known inputs and unknown class labels, and the goal is to use the available training and working data to develop a classifier that classifies the working examples with a low error rate. For results on validation of network classifiers (rather than reconciliation algorithms) in transductive settings, refer to [10] and [11]. For theory and insight on why collective classification succeeds in general settings and validation methods for it, refer to [12]. For network reconciliation, we assume that we know some network data, consisting of some node data and the links, for both networks involved in the matching, and our goal is to use that network data to match nodes as accurately as possible between the networks.

This paper presents a technique to compute probably approximately correct (PAC) bounds on the precision and recall of matching algorithms. When applied to network reconciliation algorithms, the technique does not require knowledge of how the networks were generated or rely on any specific model of network generation. The technique requires some verified matches. However, it produces valid bounds even if the verified matches are used to develop the matching algorithm, by taking advantage of a withhold-and-gap (WAG) [13] strategy – first withhold some data, train a holdout matching algorithm, validate its precision and recall on the holdout data, then train a matching algorithm on all data, and use unlabeled data to evaluate the rate of differences in matches between the two algorithms. The precision and recall for the matching algorithm trained on all data cannot be worse than the precision and recall of the holdout matching algorithm plus the rates of differences between the two algorithms, because in the worst case all differences are errors by the matching algorithm trained on all data.

The validation of precision and recall for matching algorithms is similar to the validation of classifiers over joint distributions of inputs and outputs. Both use empirical means over samples to construct bounds. However, in our setting, the sets of items to match are known, and we use empirical means over samples drawn without replacement

- *Ya Le: yle@stanford.edu, Stanford University, Dept. of Statistics*
- *Eric Bax (baxhome@yahoo.com), Nicola Barbieri, David Garcia Soriano, and Jitesh Mehta: Yahoo*
- *James Li: Facebook.*

from a finite population to bound means over the finite population. In most validation settings, samples from a joint input-output distribution are drawn with replacement and used to bound means over the input-output distribution. So, as the basis for our bounds on precision and recall for matching, we need slightly different empirical-mean bounds than those typically used in machine learning.

This paper is organized as follows. Section 2 reviews the empirical-mean bounds that we need to validate matching in our setting. Section 3 presents techniques to validate precision and recall for holdout matching algorithms. Section 4 presents techniques to extend the bounds on precision and recall to matching algorithms developed using all available data. Section 5 extends the bounds on precision and recall to other types of error bounds for matching and to some applications of matching beyond network reconciliation. Section 6 explains how to sample validation data from a finite population in our transductive setting. Section 7 offers some numerical results. Section 8 concludes with some directions for future research.

2 Bounds for Population Means

The validation methods in later sections use means over samples drawn from a finite population to estimate and bound population means. The mean bounds most commonly used in machine learning, such as [14], [15], [16], and [17], use means over samples from a distribution to bound means over the distribution. There are similar methods for population means; we review them here.

Let X be a size-n finite population, and let S be a size-s random sample drawn uniformly at random without replacement from X. Let f denote a real-valued function defined on X, i.e., $f : X \to \mathbb{R}$. Assume that f is bounded: there exist a and b such that $a \leq f(x) \leq b, \forall x \in X$.

The goal is to estimate the population mean:

$$\mu = \mathbb{E}_{x \sim \mathcal{U}(X)} \{f(x)\} = \frac{1}{n} \sum_{x \in X} f(x) \tag{1}$$

and to compute PAC bounds for it, using the sample S. (We use $\mathcal{U}()$ to denote a uniform distribution.) The sample mean:

$$\hat{\mu} = \mathbb{E}_{x \sim \mathcal{U}(S)} \{f(x)\} = \frac{1}{s} \sum_{x \in S} f(x) \tag{2}$$

is an unbiased estimator for μ. Let $p^+(X, S, f, a, b, \delta)$ and $p^-(X, S, f, a, b, \delta)$ denote PAC upper and lower bounds on μ, with bound failure probability at most δ:

$$Pr\{\mu > p^+(X, S, f, a, b, \delta)\} \leq \delta, \tag{3}$$
$$Pr\{\mu < p^-(X, S, f, a, b, \delta)\} \leq \delta. \tag{4}$$

We will outline some methods to compute these bounds.

2.1 Hoeffding's Inequality

Hoeffding [14] and Chvátal [18] show that the well-known Hoeffding inequality for distribution means also applies to population means:

$$p^+(X, S, f, a, b, \delta) = \hat{\mu} + (b - a)\sqrt{\frac{\ln 1/\delta}{2s}}, \tag{5}$$

$$p^-(X, S, f, a, b, \delta) = \hat{\mu} - (b - a)\sqrt{\frac{\ln 1/\delta}{2s}}. \tag{6}$$

Hoeffding and Chvátal also offer tighter bounds if the variance of f is known to be small. (The bounds above are based on a worst-case assumption about the variance of f.)

2.2 Empirical Bernstein-Serfling Inequality

Audibert [16] developed, and Maurer and Pontil [17] improved, empirical Bernstein bounds, named for a bound by Bernstein [19]. Empirical Bernstein bounds use the variance of f over the sample to allow the use of tighter versions of traditional bounds if the standard deviation of f is small compared to its range, which is common for error rates, since error rates tend to be small in cases of interest. Bardenet and Maillard [20] have developed similar bounds, called empirical Bernstein-Serfling bounds, for population means:

$$p^+(X, S, f, a, b, \delta) \tag{7}$$

$$= \hat{\mu} + \hat{\sigma}_s \sqrt{\frac{2\rho_s \log(5/\delta)}{s}} + \frac{\kappa(b - a)\log(5/\delta)}{s} \tag{8}$$

and

$$p^-(X, S, f, a, b, \delta) \tag{9}$$

$$= \hat{\mu} - \hat{\sigma}_s \sqrt{\frac{2\rho_s \log(5/\delta)}{s}} - \frac{\kappa(b - a)\log(5/\delta)}{s}, \tag{10}$$

where

$$\rho_s = \begin{cases} 1 - (s-1)/n & \text{if } s \leq n/2 \\ (1 - s/n)(1 + 1/n) & \text{if } s > n/2 \end{cases}, \tag{11}$$

$\hat{\sigma}_s^2 = \sum_{i,j=1}^{s} (x_i - x_j)^2/(2s^2)$, and $\kappa = \frac{7}{3} + \frac{3}{\sqrt{2}}$.

2.3 Direct Computation

If f can only take values 0 or 1, then we can use direct computation to produce tight bounds. For means over distributions, we can use binomial tail inversion, as outlined by Hoel [21] and Langford [22]. For population means, we can use hypergeometric tail inversion, from Chvátal [18], as follows.

If f can only take value 0 or 1, then $s\hat{\mu} = \sum_{i=1}^{s} f(x_i)$ follows a hypergeometric distribution. Let $m = n\mu = \sum_{x \in X} f(x)$. Then

$$Pr\{s\hat{\mu} = k\} = \binom{m}{k}\binom{n-m}{s-k} \Big/ \binom{n}{s}. \tag{12}$$

The hypergeometric tail distribution is

$$H^+(m, n, s, k) = \sum_{j=k}^{s} Pr\{s\hat{\mu} = j\} \tag{13}$$

$$= \sum_{j=k}^{s} \binom{m}{j}\binom{n-m}{s-j} \Big/ \binom{n}{s} \tag{14}$$

and

$$H^-(m,n,s,k) = \sum_{j=0}^{k} P\{s\hat{\mu} = j\} \quad (15)$$

$$= \sum_{j=0}^{k} \binom{m}{j}\binom{n-m}{s-j} \Big/ \binom{n}{s}. \quad (16)$$

So we have PAC bounds:

$$p^+(X,S,f,a,b,\delta) = \max\{p : H^-(np,n,s,s\hat{\mu}) \geq \delta\} \quad (17)$$

and

$$p^-(X,S,f,a,b,\delta) = \min\{p : H^+(np,n,s,s\hat{\mu}) \geq \delta\}, \quad (18)$$

which are tight bounds for μ, except for any rounding or approximation errors from the computation. Use either Loader's method [23] or Stirling's approximation to compute H^- and H^+.

3 Validation for Holdout Algorithms

In this section we present PAC bounds on precision and recall for holdout algorithms, which are developed while withholding a set of validation data. Let X and Y be the sets for which we match elements. (For example, X may be the set of nodes in one network, and Y may be the nodes in another.) The validation data is a sample from X for which all true matches are known and algorithm matches are computed. (For example, in matching nodes between social networks, the validation data is a set of nodes in one network for which any true matches in the other network are known.) In the next section, we extend the bounds to algorithms developed using all available data.

Let T be the set of all true matches in $X \times Y$. For $x \in X$, let $T(x)$ be the set of actual matches that include x. Let $H(x)$ be the set of x-y pairs that the holdout algorithm computes as matches that include x. Let (x,y) be a pair selected uniformly at random from $X \times Y$.

Define single-node precision:

$$p_H(x) = \frac{H(x) \cap T(x)}{H(x)} \quad (19)$$

$$= Pr\{(x,y) \in T(x) | (x,y) \in H(x)\}. \quad (20)$$

If the algorithm computes one match for x and there is at most one true match, then $p_H(x)$ is one if the algorithm match is a true match and zero otherwise. In general, $p_H(x)$ is the fraction of algorithm matches for x that are true matches. If the algorithm computes no matches for x, then $p_H(x)$ is undefined. So let

$$X_H = \{x \in X | H(x) \neq \emptyset\}, \quad (21)$$

and define holdout precision:

$$P_H = \mathbb{E}_{x \sim \mathcal{U}(X_H)}\{p_H(x)\}. \quad (22)$$

If the algorithm computes a match for each x and there is a true match for each x, then P_H is the precision as usually defined in information-retrieval settings.

Similarly, define single-node recall:

$$r_H(x) = \frac{H(x) \cap T(x)}{T(x)} \quad (23)$$

$$= Pr\{(x,y) \in H(x) | (x,y) \in T(x)\}. \quad (24)$$

If there are no actual matches for x, then $r_H(x)$ is undefined. So let

$$X' = \{x \in X | T(x) \neq \emptyset\}, \quad (25)$$

and define holdout recall to be

$$R_H = \mathbb{E}_{x \sim \mathcal{U}(X')}\{r_H(x)\}. \quad (26)$$

Let S_X be a sample drawn from X, uniformly at random without replacement. Suppose we know all true matches for each $x \in S_X$ and we compute their algorithm matches. Then we can use the sample mean for $p_H(x)$ over $x \in S_X \cap X_H$ and for $r_H(x)$ over $s \in S_X \cap X'$ to bound holdout precision and recall, respectively.

Theorem 1 (Holdout Precision and Recall). For any bound failure probability $\delta > 0$, with probability at least $1-\delta$:

$$P_H \geq p^-(X_H, S_X \cap X_H, p_H, 0, 1, \delta). \quad (27)$$

Also, with probability $1-\delta$:

$$R_H \geq p^-(X', S_X \cap X', r_H, 0, 1, \delta). \quad (28)$$

Proof: Since the elements of S_X are selected uniformly at random without replacement from X, $S_X \cap X_H$ has the same distribution as selecting a size-$|S_X \cap X_H|$ sample from X_H uniformly at random without replacement. Similarly, $S_X \cap X'$ has the same distribution as selecting a size-$|S_X \cap X'|$ sample from X' uniformly at random without replacement. □

Theorem 1 is valid if we continue to draw samples from S_X until $S_X \cap X_H$ or $S_X \cap X'$ reaches a desired sample size. If there is a method to sample directly from X_H or X', then those samples also produce valid bounds.

4 Validation for Complete Algorithms

In this section, we bound precision and recall for matching algorithms developed using all available data, which we call complete algorithms. To do this, we bound the rate of disagreement between a holdout algorithm and the complete algorithm. In the worst case, every disagreement is a mismatch by the complete algorithm. So the mismatch rate for the complete algorithm is no greater than the rate for the holdout algorithm plus the rate of disagreement.

Let C be the set of identified matches returned by the complete classfier, which can be developed using all available data, including the validation data. Let S'_X be a sample independent of S_X and drawn uniformly at random without replacement from X. Assume that $C(x)$ and $H(x)$ are computed for $x \in S_X$ and $x \in S'_X$. Define $p(x)$, X_C, P, $r(x)$, and R as complete versions of the holdout-oriented definitions in Section 3:

$$p(x) = \frac{C(x) \cap T(x)}{C(x)}, \quad (29)$$

$$X_C = \{x \in X | C(x) \neq \emptyset\}, \quad (30)$$

$$P = \mathbb{E}_{x \sim \mathcal{U}(X_C)}\{p(x)\}, \quad (31)$$

$$r(x) = \frac{C(x) \cap T(x)}{T(x)}, \quad (32)$$

$$R = \mathbb{E}_{x \sim \mathcal{U}(X')}\{r(x)\}. \quad (33)$$

Here are two remarks before presenting the results:

1) We use an extra sample, S'_X, to validate rates of disagreement between the holdout and complete algorithms. For this sample, we only need to know the holdout algorithm matches and the complete algorithm matches – not the true matches. Therefore, collecting data for S'_X is usually cheaper than collecting validation data for S_X, which includes true matches. So we can usually make the size of sample S'_X very large, producing a tight bound on the rate of disagreement between the holdout and complete algorithms.
2) For simplicity of notation, we use the same bound failure probability δ for each term in our results. But, in practice, we can use different bound failure probabilities for different terms.

Theorem 2 (Complete Recall). For any bound failure probability $\delta > 0$, with probability at least $1 - 3\delta$,

$$R \geq p^- \left(X', S_X \cap X', r_H, 0, 1, \delta\right) \tag{34}$$

$$- \frac{p^+ \left(X, S'_X, d_r, 0, 1, \delta\right)}{p^- \left(X, S_X, \mathbf{1}_{X'}, 0, 1, \delta\right)}, \tag{35}$$

where $d_r(x) = \mathbf{1}_{\{H(x) - C(x) \neq \emptyset\}}$.

Proof: Since S_X is a sample selected uniformly at random without replacement from X, by Theorem 1, the recall of the holdout algorithm satisfies

$$Pr\left\{R_H \geq p^- \left(X', S_X \cap X', r_H, 0, 1, \delta\right)\right\} \geq 1 - \delta. \tag{36}$$

The difference in recall between the complete and holdout algorithms is

$$\begin{aligned}
& R - R_H && (37)\\
= \; & \frac{1}{|X'|} \sum_{x \in X'} \frac{|C(x) \cap T(x)| - |H(x) \cap T(x)|}{|T(x)|} && (38)\\
\geq \; & -\frac{1}{|X'|} \sum_{x \in X'} \frac{|(H(x) - C(x)) \cap T(x)|}{|T(x)|} && (39)\\
\geq \; & -\frac{1}{|X'|} \sum_{x \in X'} \mathbf{1}_{\{H(x) - C(x) \neq \emptyset\}} && (40)\\
\geq \; & -\frac{|X|}{|X'|} \frac{1}{|X|} \sum_{x \in X} d_r(x). && (41)
\end{aligned}$$

Since S'_X is independent of S_X, S'_X is also a sample drawn uniformly at random without replacement from X conditional on S_X. Note that $C(x)$ is a fixed function conditional on S_X since we assume that the training samples of the holdout algorithm are fixed, therefore we can use S'_X to validate $\frac{1}{|X|} \sum_{x \in X} d_r(x)$ conditional on S_X:

$$Pr\left\{\frac{1}{|X|} \sum_{x \in X} d_r(x) \leq p^+ \left(X, S'_X, d_r, 0, 1, \delta\right) | S_X\right\} > 1 - \delta. \tag{42}$$

We derive the unconditional bound by integrating out S_X:

$$Pr\left\{\frac{1}{|X|} \sum_{x \in X} d_r(x) \leq p^+ \left(X, S'_X, d_r, 0, 1, \delta\right)\right\} > 1 - \delta. \tag{43}$$

Since $\frac{|X'|}{|X|} = \frac{1}{|X|} \sum_{x \in X} \mathbf{1}_{X'}(x)$, we can apply the bounds from Section 2:

$$Pr\left\{\frac{|X'|}{|X|} \geq p^- \left(X, S_X, \mathbf{1}_{X'}, 0, 1, \delta\right)\right\} > 1 - \delta. \tag{44}$$

Combining (36), (43), and (44) by the union bound we get the result stated in the theorem. □

Theorem 3 (Complete Precision). For any bound failure probability $\delta > 0$, with probability at least $1 - 4\delta$,

$$P \geq \tag{45}$$

$$\left[p^- (X, S'_X, \mathbf{1}_{X_H}, 0, 1, \delta) p^- (X_H, S_X \cap X_H, p_H, 0, 1, \delta)\right. \tag{46}$$

$$\left. - p^+ \left(X, S'_X, d_p, 0, 3, \delta\right)\right] \tag{47}$$

$$/ p^+ (X, S'_X, \mathbf{1}_{X_C}, 0, 1, \delta), \tag{48}$$

where

$$d_p(x) = \mathbf{1}_{X_C \cap X_H}(x) \mathbf{1}_{\{C(x) \neq H(x)\}} \tag{49}$$

$$\left(1 + \frac{|H(x) - C(x)|}{|C(x)|}\right) + \mathbf{1}_{X_H - X_C}(x). \tag{50}$$

Proof: Note that

$$\begin{aligned}
P \; = \; & \frac{1}{|X_C|} \sum_{x \in X_C} \frac{|C(x) \cap T(x)|}{|C(x)|} && (51)\\
= \; & \frac{1}{|X_C|} \left(\sum_{x \in X_C} \frac{|C(x) \cap T(x)|}{|C(x)|} - \sum_{x \in X_H} \frac{|H(x) \cap T(x)|}{|H(x)|}\right)\\
& + \frac{|X_H|}{|X_C|} \frac{1}{|X_H|} \sum_{x \in X_H} \frac{|H(x) \cap T(x)|}{|H(x)|}. && (52)
\end{aligned}$$

The first term can be decomposed into three parts

$$\begin{aligned}
& \sum_{x \in X_C} \frac{|C(x) \cap T(x)|}{|C(x)|} - \sum_{x \in X_H} \frac{|H(x) \cap T(x)|}{|H(x)|} && (53)\\
= \; & \sum_{x \in X_C \cap X_H} \left(\frac{|C(x) \cap T(x)|}{|C(x)|} - \frac{|H(x) \cap T(x)|}{|H(x)|}\right)\\
& + \sum_{x \in X_C - X_H} \frac{|C(x) \cap T(x)|}{|C(x)|}\\
& - \sum_{x \in X_H - X_C} \frac{|H(x) \cap T(x)|}{|H(x)|} && (54)\\
= \; & \mathrm{I} + \mathrm{II} + \mathrm{III}. && (55)
\end{aligned}$$

Since $|X_C - X_H|$ and $|X_H - X_C|$ are usually small, we can simply use 0 and $-|X_H - X_C|$ as the lower bound for II and III. For I, we have

$$\begin{aligned}
\mathrm{I} \; \geq \; & \sum_{\substack{x \in X_C \cap X_H \\ C(x) \neq H(x)}} \left(\frac{|H(x) \cap T(x)|}{|C(x)|} - \frac{|H(x) \cap T(x)|}{|H(x)|}\right.\\
& \left. - \frac{|(H(x) - C(x)) \cap T(x)|}{|C(x)|}\right) && (56)\\
\geq \; & \sum_{\substack{x \in X_C \cap X_H \\ C(x) \neq H(x)}} \left(-1 - \frac{|H(x) - C(x)|}{|C(x)|}\right). && (57)
\end{aligned}$$

Therefore,

$$\begin{aligned} & P \\ \geq & -\frac{1}{|X_C|} \\ & \left[\sum_{\substack{x\in X_C\cap X_H \\ C(x)\neq H(x)}} \left(1+\frac{|H(x)-C(x)|}{|C(x)|}\right)+|X_H-X_C|\right] \\ & +\frac{|X_H|}{|X_C|}\frac{1}{|X_H|}\sum_{x\in X_H}\frac{|H(x)\cap T(x)|}{|H(x)|} \qquad (58) \\ = & \frac{|X|}{|X_C|}\Bigg\{-\frac{1}{|X|}\sum_{x\in X}\Big[\mathbf{1}_{X_C\cap X_H}(x)\mathbf{1}_{\{C(x)\neq H(x)\}} \\ & \left(1+\frac{|H(x)-C(x)|}{|C(x)|}\right)+\mathbf{1}_{X_H-X_C}(x)\Big] \\ & +\frac{|X_H|}{|X|}\frac{1}{|X_H|}\sum_{x\in X_H}\frac{|H(x)\cap T(x)|}{|H(x)|}\Bigg\} \qquad (59) \end{aligned}$$

Note that

$$P_H=\frac{1}{|X_H|}\sum_{x\in X_H}\frac{|H(x)\cap T(x)|}{|H(x)|}. \qquad (60)$$

By Theorem 1, we have a lower bound for the precision of the holdout algorithm with probability at least 1-δ,

$$P_H \geq p^-\left(X_H, S_X\cap X_H, p_H, 0, 1, \delta\right). \qquad (61)$$

Since S'_X is independent of S_X, S'_X is also a sample selected uniformly at random without replacement from X conditional on S_X. Note that $C(x)$ is a fixed function conditional on S_X, therefore we can use S'_X to validate $\frac{1}{|X|}\sum_{x\in X} d_p(x)$ conditional on S_X, where $0\leq d_p(x)\leq 3$:

$$Pr\left\{\frac{1}{|X|}\sum_{x\in X} d_p(x)\leq p^+\left(X, S'_X, d_p, 0, 3, \delta\right)|S_X\right\}>1-\delta. \qquad (62)$$

We derive the unconditional bound by integrating out S_X:

$$Pr\left\{\frac{1}{|X|}\sum_{x\in X} d_p(x)\leq p^+\left(X, S'_X, d_p, 0, 3, \delta\right)\right\}>1-\delta. \qquad (63)$$

Since

$$\frac{|X_C|}{|X|}=\frac{1}{|X|}\sum_{x\in X}\mathbf{1}_{X_C}(x) \qquad (64)$$

and

$$\frac{|X_H|}{|X|}=\frac{1}{|X|}\sum_{x\in X}\mathbf{1}_{X_H}(x), \qquad (65)$$

using the bounds from Section 2, we have PAC bounds

$$Pr\left\{\frac{|X_C|}{|X|}\leq p^+\left(X, S'_X, \mathbf{1}_{X_C}, 0, 1, \delta\right)\right\} > 1-\delta, \qquad (66)$$

$$Pr\left\{\frac{|X_H|}{|X|}\geq p^-\left(X, S'_X, \mathbf{1}_{X_H}, 0, 1, \delta\right)\right\} > 1-\delta. \qquad (67)$$

Combining (61), (63), (66) and (67) by the union bound we get the result stated in the theorem. □

5 Extensions

This section outlines a few ways to extend the results from the previous sections. One extension provides different confidence levels for matches identified by different algorithms or identified by a single algorithm that provides match scores as well as matches. Other extensions bound matching error rate, instead of precision or recall, and validate matching between data fields, instead of nodes in different networks.

5.1 Confidence Levels for Different Types of Matches

The methods described in previous sections validate a single matching algorithm. In some applications, multiple matching algorithms are applied, producing multiple sets of algorithm matches. Similarly, some network reconciliation algorithms return a score for each match. The score is designed to indicate a level of similarity between the nodes in a matched pair. In this case, the range of scores may be partitioned, and scores in different parts of the range can be placed in different sets of algorithm matches.

Use union bounds to produce simultaneous PAC bounds for precision or recall over the different sets of algorithm matches. For example, if there are two sets of algorithm matches, and we validate the recall of one to be at least 80% and the other to be at least 90%, each with bound failure probability at most 2.5%, then those recall rates both hold with probability at least 100% - $2(2.5\%)$ = 95%. Simultaneous bounds support different levels of confidence for different sets of identified matches, which may affect how we choose to handle the matches in applications. For example, in an application to merge contact information based on matching user profiles of contacts across social networks, if we have only 80% precision for a set of identified matches, then we may choose to ask the user to verify matches from that set. If we have 99% precision from a different set, we may choose to merge those matches automatically, without asking the user for confirmation.

5.2 Validation of Error Rates

The validation methods for precision and recall can be easily extended to other measures of accuracy for matching algorithms. For example, for complete algorithms, we can define single-node error for node $x\in X$ as

$$w(x)=\mathbf{1}_{C(x)\neq T(x)}, \qquad (68)$$

meaning that the algorithm has an error on a node if it either fails to identify all matches for the node or identifies any false matches for it. The error rate for the algorithm can be defined as the average of single-node errors over all nodes in X. As we did for precision and recall rates, we can validate the error rate by using a sample of nodes with verified actual matches to bound the error rate for a holdout algorithm and using a sample of nodes to validate the rate of disagreement between the holdout algorithm and the algorithm developed on all available data.

5.3 Entity Resolution

In some applications, we are given a set of data fields – for example names, phone numbers, email addresses, and locations – and some information about links between them – for example that they are in the same entry of a contact list or that they occur together in a document. The goal is to match pairs of data fields that refer to the same entity or to aggregate data fields by entity. This problem is sometimes called entity resolution [1]. We can extend our validation methods to entity resolution algorithms for field matching or aggregation, as follows.

An example of matching pairs of fields is an application that connects an email address to a phone number in case someone wants to reply to an email by phone. For matching pairs of fields, we can set $X = Y$, so that we are matching nodes (fields, in our case) from a set to other nodes in the same set. Do not include pairs consisting of the same node twice in the sets of actual matches T or algorithm matches C, since we want to validate the matching of different fields to each other.

An example of field aggregation is an application that extracts fields from free text and structured data, then aggregates those fields for each of a set of people or organizations. For these matching algorithms, we can use X as the set of entities and Y as the set of fields. In general, each $x \in X$ will match to multiple $y \in Y$. In many applications, it is also possible for multiple $x \in X$ to match a single $y \in Y$. For example, multiple people may share a single home phone number and address.

6 Sampling Procedure for Validation Over a Finite Population

Our validation methods use samples S_X, drawn uniformly at random without replacement from X. In some applications, we wish to use only a subset of our labeled data as validation data, so we can also use some labeled data to develop the holdout algorithm. If we simply partition a sample into a validation set to be withheld and a training set to develop the holdout algorithm, then the holdout algorithm is not developed independently of the selection of the validation subsample, as required for our bounds to be valid, because the training subsample specifically excludes the validation subsample. Instead, we must allow a chance for some intersection between subsamples so that they are distributed as if the training and validation samples were drawn separately, and each draw was uniformly at random without replacement from X.

Let $L \subseteq X$ be a sample of X drawn uniformly at random without replacement for which we have verified the true matches. Select a validation subsample size s and a training subsample size t such that $s + t = |L|$. Generate holdout development (training) subsample D (with $|D| = t$) and validation subsample S (with $|S| = s$) as follows:

1) Select t samples uniformly at random without replacement from L. These samples are D.
2) Generate intersection size i from a hypergeometric distribution with population size $|X|$, number of success states in the population t and number of draws s. (So $Pr\{i\} = \frac{\binom{t}{i}\binom{|X|-t}{s-i}}{\binom{|X|}{s}}$.)

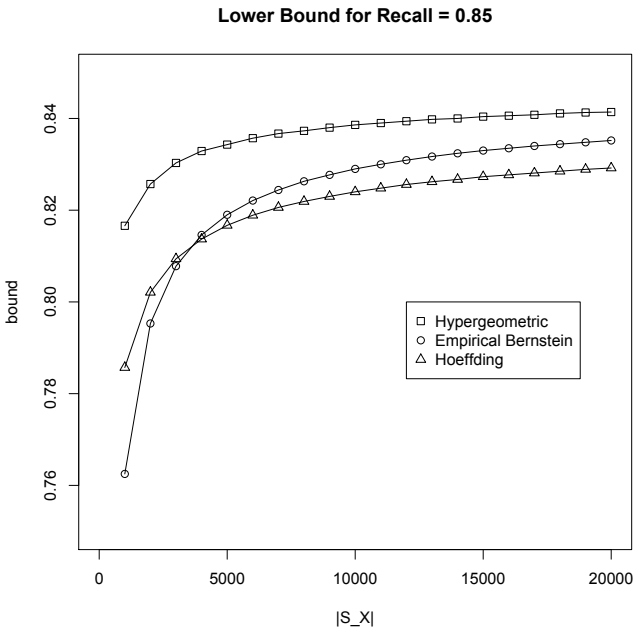

Fig. 1. Lower bounds for complete recall using Theorem 2.

3) Select i samples uniformly at random without replacement from D.
4) Select $s - i$ samples uniformly at random without replacement from $L - D$.
5) The validation subsample S consists of the samples from Steps 3 and 4.

Let D' and S' denote two independent samples of size t and s, respectively, drawn uniformly at random without replacement from the population X. Then the subsamples D and S generated by the above procedure have the same distribution as D' and S'. This is because the size of the intersection between D and S has the same distribution as that between D' and S', and the samples drawn in Steps 1, 3, and 4 have the same distribution as those drawn uniformly at random without replacement from the population. See the appendix for a rigorous proof.

7 Numerical Results

This section presents some numerical results for complete recall for typical parameter values for matching users between social networks. This gives a general sense of bound effectiveness. It also allows us to compare the different methods for bounding actual means based on empirical means: Hoeffding, empirical Bernstein, and hypergeometric (direct). Let $|X| = 120$ million. Assume that the fraction of $x \in X$ that have a match in Y is $\frac{2}{3}$, so $|X'| = 80$ million. Assume the rate of disagreement between holdout and complete algorithms, d_r, is 0.001. Let $|S'_X|$ = 100,000. (These are samples used to assess the rate of disagreement; their actual matches need not be verified.) Let $\delta = \frac{0.05}{3}$, so that we have a 95% bound success probability.

Figure 1 shows lower bounds for complete recall using Theorem 2 if the holdout matching algorithm has 85% recall ($r_H = 0.85$). Using hypergeometric bounds for

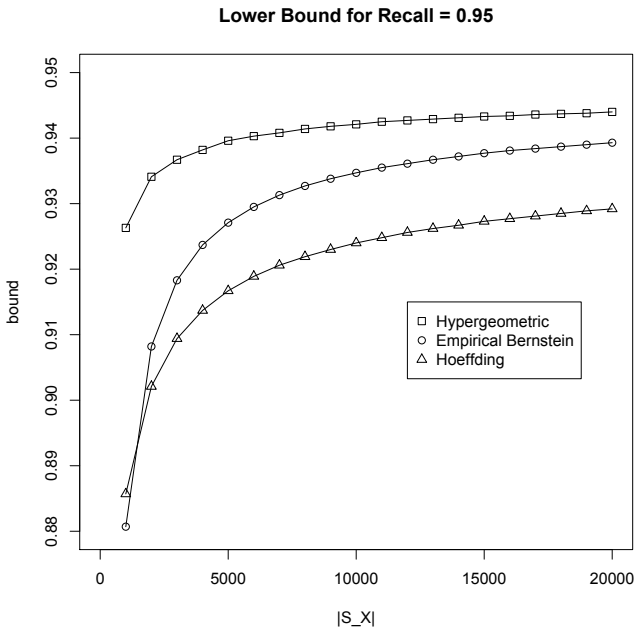

Fig. 2. Lower bounds for complete recall using Theorem 2.

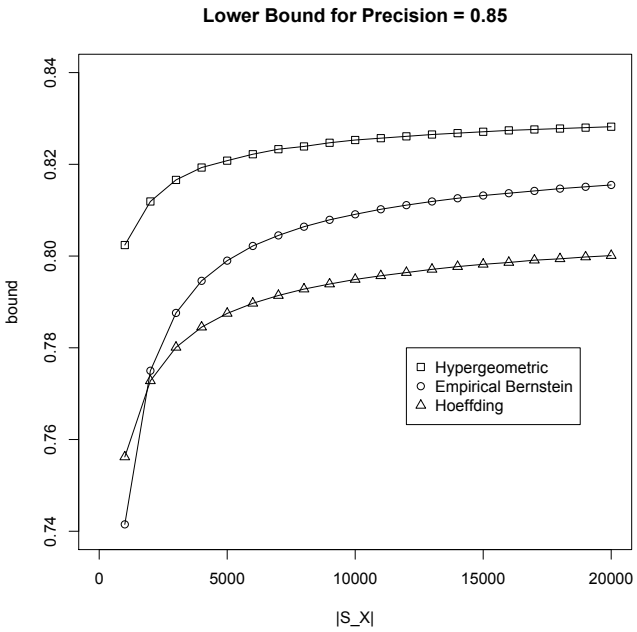

Fig. 3. Lower bounds for complete precision using Theorem 3.

$p^+()$ and $p^-()$, the lower bound is above 81% even using $|S_X| = 1000$: a thousand samples from X with verified matches. Hypergeometric bounds outperform the others for all numbers of samples, since they are sharp bounds (up to machine rounding errors and any error from using Stirling's approximation for factorials.) Empirical Bernstein bounds are stronger than Hoeffding bounds for 5000 or more verified matches.

Figure 2 shows lower bounds for complete recall if the holdout matching algorithm has 95% recall. With a more accurate matching algorithm, empirical Bernstein bounds begin to outperform Hoeffding bounds even for 2000 verified matches. Even with only 1000 samples with verified matches, hypergeometric bounds produce a lower bound on complete recall of over 92%.

Figures 3 and 4 show lower bounds for complete precision based on Theorem 3. We use the same parameter values as for the recall results in Figures 1 and 2, except $\delta = \frac{0.05}{4}$, to keep a five percent bound failure probability. We assume there is at most one holdout algorithm match and at most one complete algorithm match for each $x \in X$. We assume both algorithms match $\frac{2}{3}$ of the elements of X, and $d_p = 0.003$. The holdout precision rate p_H is 0.85 for Figure 3 and 0.95 for Figure 4.

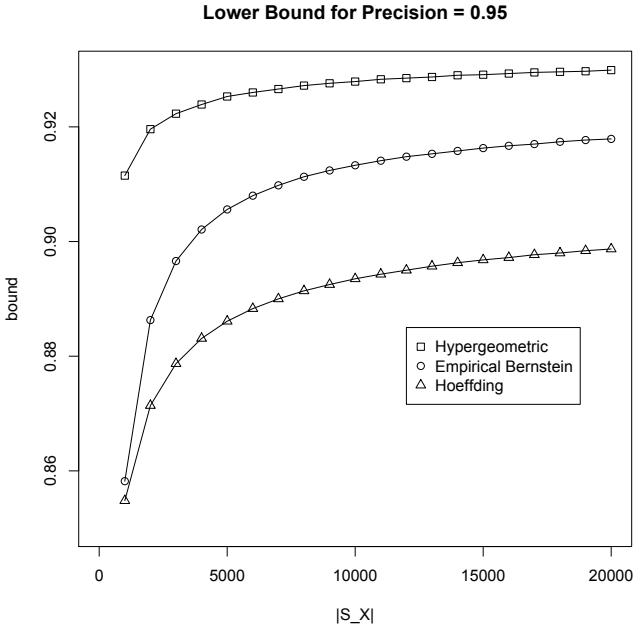

Fig. 4. Lower bounds for complete precision using Theorem 3.

8 Discussion

This paper introduces methods to validate precision and recall rates for network reconciliation algorithms. The methods apply regardless of the process that generated the networks' link structure and node data. In fact, the methods apply to matching algorithms in general, not just those that match nodes between networks. The methods include validation of matching algorithms that use all available data, so data used for validation may be used for matching as well.

One direction for future research is to use the validation methods developed in this paper to guide network reconciliation processes. For example, some iterative matching algorithms begin with a seed set of matches and use matches from previous iterations to infer new matches in successive iterations, including methods from [6], [4], [7], and [5]. In each iteration, validated precision and recall rates may be useful to determine which pairs to adopt as new matches and how much to rely on them to infer future matches. Pedarsani et al. [7] use Bayesian confidence measures in this

way, under the assumption that the features used for matching have independent distributions. The validation methods developed in this paper could be used in a similar way when the feature distributions are correlated or unknown.

Finally, it would be interesting to develop validation methods for network reconciliation in non-transductive settings, where the goal is to assess precision and recall for nodes to be added to the network in the future, whose data and links are not yet known. For these settings, we may need to assume that we know how the network will grow, or at least have a probabilistic model for network growth. Alternatively, we may be able to develop validation methods based on the assumption that soon-to-be-added nodes will be generated by the same distribution as recently-added nodes, even if we do not know the distribution, as in work on algorithm validation using cohorts [11].

REFERENCES

[1] I. Bhattacharya and L. Getoor, "Collective entity resolution in relational data," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 1, p. 5, 2007.

[2] Y. Tian and J. M. Patel, "Tale: A tool for approximate large graph matching," in *IEEE 24th International Conference on Data Engineering, 2008. ICDE 2008*, 2008, pp. 963–972.

[3] M. Zaslavskiy, F. Bach, and J.-P. Vert, "Path following algorithm for the graph matching problem," *CoRR*, vol. abs/0801.3654, 2008.

[4] N. Korula and S. Lattanzi, "An efficient reconciliation algorithm for social networks," *Proceedings of the 40th International Conference on Very large Databases (VLDB)*, pp. 377–388, 2014.

[5] L. Yartseva and M. Grossglauser, "On the performance of percolation graph matching," *Conference on Online Social Networks (COSN)*, pp. 119–129, 2013.

[6] A. Nunes, P. Calad, and B. Martins, "Resolving user identities over social networks through supervised learning and rich similarity features," in *Proceedings of the 27th Annual ACM Symposium on Applied Computing*, ser. SAC '12. New York, NY, USA: ACM, 2012, pp. 728–729. [Online]. Available: http://doi.acm.org/10.1145/2245276.2245413

[7] P. Pedarsani, D. R. Figueiredo, and M. Grossglauser, "A Bayesian method for matching two similar graphs without seeds," in *Fifty-first Annual Allerton Conference*, 2013, pp. 1598–1607.

[8] P. Foggia, G. Percannella, and M. Vento, "Graph matching and learning in pattern recognition in the last 10 years," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 28, no. 01, p. 1450001, 2014. [Online]. Available: http://www.worldscientific.com/doi/abs/10.1142/S0218001414500013

[9] V. Vapnik, *Statistical Learning Theory*. John Wiley & Sons, 1998.

[10] J. Li, A. Sonmez, Z. Cataltepe, and E. Bax, "Validation of network classifiers," *Structural, Syntactic, and Statistical Pattern Recognition Lecture Notes in Computer Science*, vol. 7626, pp. 448–457, 2012.

[11] E. Bax, J. Li, A. Sonmez, and Z. Cataltepe, "Validating collective classification using cohorts," *NIPS Workshop: Frontiers of Network Analysis: Methods, Models, and Applications*, 2013.

[12] B. London, B. Huang, B. Taskar, and L. Getoor, "Pac-bayesian collective stability," in *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*, 2014.

[13] E. Bax and Y. Le, "Some theory for practical classifier validation," *Baylearn 2015*, 2015. [Online]. Available: http://arxiv.org/pdf/1510.02676v1.pdf

[14] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13–30, 1963.

[15] S. Boucheron, G. Lugosi, and P. Massart, *Concentration Inequalities – A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

[16] J.-Y. Audibert, "PAC-Bayesian statistical learning theory," Ph.D. dissertation, Laboratoire de Probabilities et Modeles Aleatoires, Universites Paris 6 and Paris 7, 2004. [Online]. Available: http://cermis.enpc.fr/~{}audibert/ThesePack.zip

[17] A. Maurer and M. Pontil, "Empirical Bernstein bounds and sample-variance penalization," in *COLT*, 2009.

[18] V. Chvátal, "The tail of the hypergeometric distribution," *Discrete Mathematics*, vol. 25, no. 3, pp. 285–287, 1979.

[19] S. N. Bernstein, "On certain modifications of Chebyshev's inequality," *Doklady Akademii Nauk SSSR*, vol. 17, no. 6, pp. 275–277, 1937.

[20] R. Bardenet and O.-A. Maillard, "Concentration inequalities for sampling without replacement," 2013.

[21] P. G. Hoel, *Introduction to Mathematical Statistics*. Wiley, 1954.

[22] J. Langford, "Tutorial on practical prediction theory for classification," *Journal of Machine Learning Research*, vol. 6, pp. 273–306, 2005.

[23] C. Loader, "Fast and accurate computation of binomial probabilities," 2000. [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.35.2719

APPENDIX

Proof of Subsampling Procedure from Section 6

Proof: Let $|X| = n$, $|D| = t$ and $|S| = s$. Note that $|L| = t + s$. It suffices to show that

$$Pr\{D\} = \frac{1}{\binom{n}{t}}, \tag{69}$$

$$Pr\{S|D\} = \frac{1}{\binom{n}{s}}. \tag{70}$$

Since L is selected uniformly at random without replacement from X, we have

$$Pr\{D\} = \sum_L Pr\{L\} Pr\{D|L\} \tag{71}$$

$$= \sum_{L \supseteq D} \frac{1}{\binom{n}{t+s}\binom{t+s}{t}} \tag{72}$$

$$= \frac{\binom{n-t}{s}}{\binom{n}{t+s}\binom{t+s}{t}} \tag{73}$$

$$= \frac{1}{\binom{n}{t}}. \tag{74}$$

This shows that D has the same distribution as a size-t sample selected uniformly at random without replacement from X.

The distribution of S conditional on D satisfies

$$Pr\{S|D\} = \sum_L Pr\{L|D\} Pr\{S|D, L\} \tag{75}$$

$$= \sum_{L \supseteq D \cup S} Pr\{L|D\} Pr\{i = |S \cap D|\} Pr\{S|N, D, L\} \tag{76}$$

$$= \frac{\binom{n-t-s+|S\cap D|}{|S\cap D|}\binom{t}{|S\cap D|}\binom{n-t}{s-|S\cap D|}}{\binom{n-t}{s}\binom{n}{s}\binom{t}{|S\cap D|}\binom{s}{s-|S\cap D|}} \tag{77}$$

$$\frac{1}{\binom{n}{s}}. \tag{78}$$

Thus S is independent of D and has the same distribution as a size-s sample selected uniformly at random without replacement from X. □